Stochastic Dynamic Cache Partitioning for Encrypted Content Delivery

Andrea Araldo*‡, György Dán†, Dario Rossi‡
* LRI, Université ParisSud, Orsay, France – araldo@lri.fr
† School of Electrical Engineering, KTH Royal Institute of Technology, Stockholm, Sweden – gyuri@ee.kth.se
‡ Télécom ParisTech, Paris, France – dario.rossi@telecom-paristech.fr

arXiv:1603.09490v2 [cs.NI] 27 Jun 2016

***Abstract*—In-network caching is an appealing solution to cope with the increasing bandwidth demand of video, audio and data transfer over the Internet. Nonetheless, an increasing share of content delivery services adopt encryption through HTTPS, which is not compatible with traditional ISP-managed approaches like transparent and proxy caching. This raises the need for solutions involving both Internet Service Providers (ISP) and Content Providers (CP): by design, the solution should preserve business-critical CP information (e.g., content popularity, user preferences) on the one hand, while allowing for a deeper integration of caches in the ISP architecture (e.g., in 5G femto-cells) on the other hand.**

In this paper we address this issue by considering a content-oblivious ISP-operated cache. The ISP allocates the cache storage to various content providers so as to maximize the bandwidth savings provided by the cache: the main novelty lies in the fact that, to protect business-critical information, ISPs only need to measure the aggregated miss rates of the individual CPs and do not need to be aware of the objects that are requested, as in classic caching. We propose a cache allocation algorithm based on a perturbed stochastic subgradient method, and prove that the algorithm converges close to the allocation that maximizes the overall cache hit rate. We use extensive simulations to validate the algorithm and to assess its convergence rate under stationary and non-stationary content popularity. Our results (i) testify the feasibility of content-oblivious caches and (ii) show that the proposed algorithm can achieve within 10% from the global optimum in our evaluation.

I. Introduction

It is widely known that content delivery over the Internet represents a sizeable and increasing fraction of the overall traffic demand. Furthermore most of the content, including video, is carried over HTTP connections: this evolution of the last decade was not among those forecasted for the IP hourglass model evolution [1], and is rather a choice of practical convenience. This evolution has a tremendous practical relevance, to the point that HTTP was overtly recognized and proposed [2] as the new de facto "thin waist" of the TCP/IP protocol family.

In very recent times, we are on the verge of yet another shift of the thin waist: we indeed observe that the fraction of traffic delivered through HTTPS has already passed 50% [3], and it is expected to increase, as the IETF Internet Architecture Board (IAB) recommends "protocol designers, developers, and operators to make encryption the norm for Internet traffic" [4]. Besides the IAB recommendation, Content Providers (CP) are already heavily relying on encryption to both protect the privacy of their users, as well as sensitive information (related to user preferences) of their own business.

This evolution toward an all-encrypted Internet creates a tussle between security and efficiency. Today's Internet heavily relies on middleboxes such as NATs (to combat the scarcity of IPv4 addresses) and transparent or proxy caches [5] (to relieve traffic load). However, some of these middleboxes will simply fail to operate in tomorrow's Internet with end-to-end encryption: for example, end-to-end encryption renders caching useless, since multiple transfers of the same object generate different streams if the same object is encrypted with different keys. At times where the design of the new 5G architecture strives to reduce latency, increase the available bandwidth and better handle mobility, this tradeoff is especially unfortunate, as distributed caches represent a natural way to reduce latency, reduce bandwidth usage and to cope with mobility avoiding long detours to anchor points.

This architectural evolution calls for a redesign of the current operations involving both Internet Service Providers (ISP) and Content Providers (CP): by design, the solutions should preserve business-critical CP information (e.g., content popularity, user preferences) on the one hand, while allowing for a deeper integration of caches in the ISP architecture (e.g., in 5G femto-cells) on the other hand.

In this paper we address this issue by proposing a content-oblivious algorithm that manages the storage space of an ISP cache that delivers *encrypted content*: the algorithm *dynamically partitions* the cache storage among various CPs so as to maximize the cache hit rate (hence, the bandwidth savings). The most important feature of the algorithm is that in order to protect business-critical information of the CPs, the ISP only needs to measure the *aggregated miss rates* of the individual CPs. We prove that relying on the aggregated miss rates only, our algorithm converges close to the optimal allocation that maximizes the overall cache hit rate, and provide a bound on the gap compared to the optimal allocation. Extensive simulations under realistic scenarios show the feasibility and good performance of the proposed algorithm.

The rest of the paper is organized as follows. Sec. II reviews related work. Sec. III describes the system model and presents our cache partitioning algorithm. Sec. IV provides analytic results, and Sec. V evaluates the performance of the algorithm. Sec. VI concludes the work.

II. Related Work

The problem of cache partitioning has first been considered in the context of partitioning CPU cache among competing processes [6]. Clearly, the cache workload created by CPU jobs differs very much from the characteristics of traffic requests from a user aggregate, so that our technique significantly differs in their design choices: i.e., there is no notion of privacy within the CPU that would force the scheduler to be agnostic of the instruction patterns, unlike in our case.

Related to our work are recent works in the area of in-network caching, which is an established technique that continues to receive significant interest in recent years [7]. They generally consider the cache as a whole, but for a few exceptions [8], [9] that partition a cache among different applications and video quality levels, respectively, and assume that the served content is observable by the ISP. Our work differs from these works, as we consider a cache partitioned among CPs, and the ISP can observe the aggregate miss rates of the CPs only. A recent work [10] proposes to share a cache managed by an ISP among different CPs in order to achieve fairness. Partitioning of the storage is done using a pricing scheme based on the value that each CP gives to cache space. Unlike [10], our proposed algorithm aims to maximize the cache hit rate, and does not involve payments, which may make its adoption less controversial considering the disputes among ISPs and CPs in recent years [11]. To the best of our knowledge, our work is the first to propose an effective way to maximize the efficiency of an ISP-managed cache, by partitioning it among multiple CPs, without being aware of the content that the CPs serve and without involving payments.

Closer in scope to ours are worth mentioning a set of recent works that target ISP/CDN cooperation [12], [13], [14], [15]. These not only show that ISPs have strong incentive in investing in caching to reduce the traffic on their critical paths, but also show that the other Internet actors, i.e. CPs and users, would benefit from ISP in-network caching. The game theoretical study in [12] shows that caches are inefficient when operated by CPs, since CP content placement and ISP traffic engineering are often not compatible. Solutions are proposed in [13], [15], which however require ISPs to share with the CP confidential information, such as topology, routing policies or link states, and as such are arguably highly impractical. Conversely, [14] fosters an ISP-operated cache system, but requires the ISP to be able to observe every object requested by the users, which is arguably equally impractical since CPs purposely hide this confidential information via HTTPS. In contrast with these previous works, our solution does not yield to any leaking of business critical information. Furthermore, our solution is not limited to a single CP, unlike [13].

Finally, from a technical viewpoint, our work is aligned with recent industry efforts in the Open Caching Working Group (OCWG) [16]. The mission of the OCWG is to develop standards, policies and best practices for a new layer of content caching within ISP networks, which can coexist with HTTPS and provide shared access to storage for many CPs. Our work fits the OCWG requirements and as such is, we believe, of high practical relevance.

III. Stochastic Dynamic Cache Partitioning

A. System Model and Problem Formulation

We consider a cache with a storage size of K slots (e.g., in units of MB) maintained by an operator and shared by P content providers (CPs). The operator is not aware of what content the individual slots store and only decides how to partition the slots among CPs.

We denote by $\theta_p \in \mathbb{Z}_{\geq 0}$ the number of cache slots allocated to CP p, which it can use for caching its most popular contents. We define the set of feasible cache allocation vectors

$$\Theta \triangleq \{\boldsymbol{\theta} \in \mathbb{Z}_{\geq 0}^{P} | \sum_{p=1}^{P} \theta_p \leq K\} \subset \mathbb{Z}_{\geq 0}^{P}. \tag{1}$$

We consider that the arrival of requests for content can be modeled by a stationary process, and the number of arrivals over a time interval of length T can be *bounded* by some positive constant $A(T)$. This assumption is reasonable as requests are generated by a finite customer population, and each customer can generate requests at a bounded rate in practice. Upon reception of a request for a content of the CPs that share the cache, the request can either generate a cache hit (for content stored in the CP partition at time of the request) or a cache miss (otherwise). Formally, we denote the *expected* cache miss rate (i.e., expected number of misses per time unit) of CP p when allocated θ_p slots of storage by $L_p(\theta_p)$. We make the reasonable assumption that L_p is decreasing and strictly convex on $[0 \dots K]$. This assumption corresponds to that having more storage decreases the miss intensity (in expectation) with a decreasing marginal gain, and each CP would in principle have enough content to fill the entire storage. For convenience we define the cache miss intensity vector

$$\vec{\mathbf{L}}(\boldsymbol{\theta}) \triangleq \left(L_1(\theta_1), \dots, L_P(\theta_P)\right)^T. \tag{2}$$

Finally, we define the overall expected cache miss intensity

$$L(\boldsymbol{\theta}) \triangleq \sum_{p=1}^{P} L_p(\theta_p). \tag{3}$$

Motivated by the increasing prevalence of encrypted content delivery, we assume that the operator cannot observe what content an individual request is for, but it can observe the number of content requests received by a CP and the corresponding number of cache misses.

Given a static cache partitioning $\boldsymbol{\theta} \in \Theta$, the observed number of content requests and the number of cache misses would form a stationary sequence when measured over subsequent time intervals. The objective of the operator is to find the optimal allocation $\boldsymbol{\theta}^{OPT}$ that minimizes the overall *expected cache miss intensity*, i.e.,

$$\boldsymbol{\theta}^{OPT} \in \arg\min_{\boldsymbol{\theta} \in \Theta} L(\boldsymbol{\theta}), \tag{4}$$

based on the measured cache miss intensity. In what follows we propose the Stochastic Dynamic Cache Partitioning (SDCP) algorithm that iteratively approximates the optimal allocation.

Table I
FREQUENTLY USED NOTATION (WITH PLACE OF DEFINITION)

P	Number of content providers (CP)
K	Available cache slots
K'	Allocated cache slots (6)
$\boldsymbol{\theta}$	Cache configuration
Θ	Set of feasible cache allocations (1)
$\mathcal{C}$	Set of allowed virtual cache allocations (5)
$L(\boldsymbol{\theta})$	Expected cache miss intensity (3)
$\vec{\mathbf{L}}$	Miss intensity vector (2)
$\bar{L}(\boldsymbol{\theta})$	Interpolant of the miss intensity (Lemma 8)
$\boldsymbol{\theta}^*$	Unique minimizer of $\bar{L}$ (Lemma 8)
$\mathbf{D}^{(k)}$	Perturbation vector (8)
T	Time slot length
$\boldsymbol{\varphi}$	Euclidean projection (10)
$\hat{\mathbf{g}}^{(k)}$	Stochastic subgradient (line 11 of Alg. 1)
$\bar{\mathbf{g}}(\boldsymbol{\theta})$	Subgradient of $\bar{L}$ (21)

Algorithm 1: Stochastic Dynamic Cache Partitioning Algorithm

1 Choose an initial allocation $\boldsymbol{\theta}_0 \in \mathcal{C} \cap \mathbb{R}^P_{\geq 0}$
2 **for** $k = 0$*; ;* $k++$ **do**
3 Generate $\mathbf{D}^{(k)}$
4 ${}^{+}\boldsymbol{\theta}^{(k)} = \boldsymbol{\Gamma}(\boldsymbol{\theta}^{(k)}) + \frac{1}{2}\mathbf{D}^{(k)}$
5 ${}^{-}\boldsymbol{\theta}^{(k)} = \boldsymbol{\Gamma}(\boldsymbol{\theta}^{(k)}) - \frac{1}{2}\mathbf{D}^{(k)}$
6 Set the configuration to ${}^{+}\boldsymbol{\theta}^{(k)}$ for time $T/2$
7 Measure ${}^{+}\vec{\mathbf{y}}^{(k)}$
8 Set the configuration to ${}^{-}\boldsymbol{\theta}^{(k)}$ for a time $T/2$
9 Measure ${}^{-}\vec{\mathbf{y}}^{(k)}$
10 $\delta\vec{\mathbf{y}}^{(k)} = {}^{+}\vec{\mathbf{y}}^{(k)} - {}^{-}\vec{\mathbf{y}}^{(k)}$
11 $\hat{\mathbf{g}}^{(k)} = \delta\vec{\mathbf{y}}^{(k)} \circ \mathbf{D}^{(k)} - \frac{1}{P} \cdot (\delta\vec{\mathbf{y}}^{(k)^T} \cdot \mathbf{D}^{(k)})\mathbf{1}_P$
12 $\boldsymbol{\theta}^{(k+1)} = \boldsymbol{\varphi}(\boldsymbol{\theta}^{(k)} - a^{(k)}\hat{\mathbf{g}}^{(k)})$
13 **end**

B. Stochastic Dynamic Cache Partitioning (SDCP) Algorithm

The proposed SDCP algorithm is an iterative algorithm that is executed over time slots of fixed duration T. The pseudo code of the algorithm is shown in Alg. 1. For simplicity we present the algorithm assuming that P is even, but the case of an odd number of CPs can be handled by introducing a fictitious CP with zero request rate.

At time slot k the algorithm maintains a virtual cache allocation $\boldsymbol{\theta}^{(k)}$. The virtual allocation is an allocation of K' storage slots among the CPs, i.e.,

$$\boldsymbol{\theta}^{(k)} \in \mathcal{C} \triangleq \left\{ \boldsymbol{\theta} \in \mathbb{R}^p | \mathbf{1}_P^T \cdot \boldsymbol{\theta} \triangleq K' \right\}, \tag{5}$$

where

$$K' = K - P/2. \tag{6}$$

We will justify the introduction of K' and of $\mathcal{C}$ in the proof of Lemma 6.

In order to obtain from $\boldsymbol{\theta}^{(k)}$ an integral allocation that can be implemented in the cache, we define the center-point function $\boldsymbol{\Gamma} : \mathbb{R}^P \to \mathbb{R}^P$, which assigns to a point in Euclidean space the center of the hypercube containing it, i.e.,

$$\begin{aligned} \gamma(x) &\triangleq \lfloor x \rfloor + 1/2, && \forall x \in \mathbb{R}, \\ \boldsymbol{\Gamma}(\boldsymbol{\theta}) &\triangleq (\gamma(\theta_1), \ldots, \gamma(\theta_P))^T, && \forall \boldsymbol{\theta} \in \mathcal{C}, \end{aligned} \tag{7}$$

where we use $\lfloor \cdot \rfloor$ to denote the floor of a scalar or of a vector in the component-wise sense. Furthermore, we define the perturbation vector $\mathbf{D}^{(k)} = (D_1^{(k)}, \ldots, D_P^{(k)})^T$ at time slot k, which is chosen independently and uniform at random from the set of -1,$+1$ valued zero-sum vectors

$$\mathbf{D}^{(k)} \in \mathcal{Z} \triangleq \left\{ \mathbf{z} \in \{-1, 1\}^P \,\middle|\, \mathbf{z}^T \cdot \mathbf{1}_P = 0 \right\}. \tag{8}$$

Given $\boldsymbol{\Gamma}$ and $\mathbf{D}^{(k)}$ the algorithm computes two cache allocations to be implemented during time slot k,

$$\begin{aligned} {}^{+}\boldsymbol{\theta}^{(k)} &\triangleq \boldsymbol{\Gamma}(\boldsymbol{\theta}^{(k)}) + \tfrac{1}{2}\mathbf{D}^{(k)}, \\ {}^{-}\boldsymbol{\theta}^{(k)} &\triangleq \boldsymbol{\Gamma}(\boldsymbol{\theta}^{(k)}) - \tfrac{1}{2}\mathbf{D}^{(k)}. \end{aligned} \tag{9}$$

The algorithm first applies allocation ${}^{+}\boldsymbol{\theta}^{(k)}$ for $T/2$ amount of time and measures the cache miss rate ${}^{+}y_p^{(k)}$ for each provider $p = 1, \ldots, P$. It then applies allocation ${}^{-}\boldsymbol{\theta}^{(k)}$ during the remaining $T/2$ amount of time in slot k and measures the cache miss rates ${}^{-}y_p^{(k)}$. The vectors of measured cache misses ${}^{-}\vec{\mathbf{y}}^{(k)} \triangleq ({}^{-}y_1^k, \ldots, {}^{-}y_P^{(k)})^T$ and ${}^{+}\vec{\mathbf{y}}^{(k)} \triangleq ({}^{+}y_1^k, \ldots, {}^{+}y_P^{(k)})^T$ are used to compute the impact $\delta y_p^{(k)} \triangleq {}^{+}y_p^{(k)} - {}^{-}y_p^{(k)}$ of the perturbation vector on the cache miss intensity of CP p, or using the vector notation $\delta\vec{\mathbf{y}}^{(k)} \triangleq {}^{+}\vec{\mathbf{y}}^{(k)} - {}^{-}\vec{\mathbf{y}}^{(k)}$.

Based on the measured miss rates, the algorithm then computes the allocation vector $\boldsymbol{\theta}^{(k+1)}$ for the $(k+1)$-th step. Specifically, it first computes (line 11, where $\circ$ denotes the Hadamard product) the update vector $\hat{\mathbf{g}}^{(k)}$, which we show in Cor. 12 to match in expectation a subgradient of the miss-stream interpolant $\bar{L}$, defined in Lemma 8. The $(k+1)$-th allocation moves from the k-th allocation in the direction of the update vector $\hat{\mathbf{g}}^{(k)}$, opportunely scaled by a *step size* $a^{(k)} > 0$. Additionally, denoting with $\mathbb{R}_{\geq 0}$ the set of non-negative numbers, $\boldsymbol{\theta}^{(k+1)}$ is computed using the Euclidean projection $\boldsymbol{\varphi} : \mathcal{C} \to \mathcal{C} \cap \mathbb{R}^P_{\geq 0}$, defined as

$$\boldsymbol{\varphi}(\boldsymbol{\theta}) \triangleq \arg \min_{\boldsymbol{\theta}' \in \mathcal{C} \cap \mathbb{R}^P_{\geq 0}} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|. \tag{10}$$

Several remarks are worth making. First, we will show in Lemma 5 that the equation above admits a unique solution and thus the definition is consistent. Second, we will show in Lemma 6 that $\hat{\mathbf{g}}$ computed as in line 11 guarantees that the update $\boldsymbol{\theta}^{(k)} - a^{(k)}\hat{\mathbf{g}}^{(k)}$ at line 12 lies inside $\mathcal{C}$. Nonetheless, this update may have some negative components and we need to project it into $\mathcal{C} \cap \mathbb{R}_{\geq 0}$ by applying $\boldsymbol{\varphi}$, to ensure that the subsequent virtual allocation $\boldsymbol{\theta}^{(k+1)}$ is valid. Third, the *step size* $a^{(k)}$ must be chosen to satisfy $\sum_{k=1}^{\infty} a^{(k)} = \infty$ and $\sum_{k=1}^{\infty} (a^{(k)})^2 < \infty$ in order to guarantee convergence (see Theor. 13). Fourth, although the convergence of the proposed algorithm is guaranteed, for stationary content popularity, irrespectively of the choice of $a^{(k)}$ satisfying the above conditions, we point out that the step size plays an important role in determining the convergence speed, which we will numerically investigate in Sec. V.

IV. CONVERGENCE ANALYSIS OF SDCP

We first provide definitions and known results (Sec.IV-A) that are instrumental to prove important properties of the proposed algorithm: consistency (Sec.IV-B), convergence (Sec.IV-C) and a bound on the optimality gap (Sec.IV-D).

A. Preliminaries

Let us start by introducing the forward difference defined for functions on discrete sets.

Definition 1. For a function $\mathbf{F} : \mathbb{Z}^{q_1} \to \mathbb{R}^{q_2}$, $q_1, q_2 \geq 1$ the forward difference is

$$\Delta_n \mathbf{F}(\mathbf{x}) \triangleq \mathbf{F}(\mathbf{x} + n \cdot \mathbf{1}_{q_1}) - F(\mathbf{x}), \forall \mathbf{x} \in \mathbb{Z}^{q_1}, n \in \mathbb{Z} \setminus \{0\}.$$

By abuse of notation, we will simply use $\Delta\mathbf{F}(\mathbf{x})$ to denote $\Delta_1\mathbf{F}(\mathbf{x})$.

The forward difference is convenient for characterizing convexity using the following definition [17].

Definition 2. A discrete function $F : \mathbb{Z} \to \mathbb{R}$ is strictly convex iff $x \to \Delta F(x)$ is increasing.

Furthermore, for a class of functions of interest we can establish the following.

Lemma 3. *Let $F : \mathbb{Z} \to \mathbb{R}$ decreasing and strictly convex, $x \in \mathbb{Z}$ and $n \in \mathbb{Z} \setminus \{0\}$ we have*

$$\Delta_n F(x) > n\Delta F(x). \tag{11}$$

Proof:

We first show that $\forall x, y \in \mathbb{Z}$ such that $y > x$, the following holds

$$\Delta_n F(y) > \Delta_n F(x) \quad \text{if } n > 0, \tag{12}$$
$$\Delta_n F(y) < \Delta_n F(x) \quad \text{if } n < 0. \tag{13}$$

For $n > 0$ we can use Def. 2 to obtain

$$\Delta_n F(y) = \sum_{i=0}^{n-1} [F(y+i+1) - F(y+i)] = \sum_{i=0}^{n-1} \Delta F(y+i)$$
$$> \sum_{i=0}^{n-1} \Delta F(x+i) = \Delta_n F(x), \tag{14}$$

which proves (12).

For $n < 0$ algebraic manipulation of the definition of the forward difference and (14) gives

$$\Delta_n F(y) = -\Delta_{|n|} F(y - |n|) < -\Delta_{|n|} F(x - |n|) = \Delta_n F(x),$$

which proves (13). To prove (11) for $n > 0$, observe that, thanks to Def. 2, each of the n terms of the last summation in (12) is lower bounded by $\Delta F(x)$. For $n < 0$ via algebraic manipulation we obtain

$$\Delta_n F(x) = -\sum_{i=1}^{|n|} \Delta F(x-i) > -\sum_{i=1}^{|n|} \Delta F(x) = -|n| \cdot \Delta F(x),$$

which proves (11) as $|n| = -n$. ∎

Since SDCP generates virtual configurations whose components are not necessarily integer, we have to extend the discrete functions L_p to real numbers. Thanks to Theor. 2.2 of [18], we have the following existence result.

Lemma 4. *Given a discrete decreasing and strictly convex function $F : \mathbb{Z} \to \mathbb{R}$, there exists a continuous and strictly convex function $\bar{F} : \mathbb{R} \to \mathbb{R}$ that extends F, i.e., $F(x) = \bar{F}(x), \forall x \in \mathbb{Z}$. We call $\bar{F}$ the* interpolant *of F.*

Finally, we formulate an important property of the Euclidean projection $\boldsymbol{\varphi}$.

Lemma 5. *There is a unique function $\boldsymbol{\varphi}$ satisfying (10). Furthermore, $\boldsymbol{\varphi}$ satisfies*

$$\|\boldsymbol{\varphi}(\boldsymbol{\theta}) - \boldsymbol{\theta}'\| \leq \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|, \forall \boldsymbol{\theta} \in \mathcal{C}, \boldsymbol{\theta}' \in \mathcal{C} \cap \mathbb{R}^P_{\geq 0}, \tag{15}$$

i.e., $\boldsymbol{\varphi}(\boldsymbol{\theta})$ is no farther from any allocation vector than $\boldsymbol{\theta}$.

Proof: Observe that $\mathcal{C} \cap \mathbb{R}^P_{\geq 0}$ is a simplex, and thus closed and convex. Hence, the Euclidean projection $\boldsymbol{\varphi}$ is the unique solution of (10) [19]. Furthermore, the Euclidean projection is non-expansive (see, e.g., Fact 1.5 in [20]), i.e., for $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathcal{C}$ it satisfies $\|\boldsymbol{\varphi}(\boldsymbol{\theta}) - \boldsymbol{\varphi}(\boldsymbol{\theta}')\| \leq \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|$. Observing that if $\boldsymbol{\theta}' \in \mathcal{C} \cap \mathbb{R}^P_{\geq 0}$ then $\boldsymbol{\varphi}(\boldsymbol{\theta}') = \boldsymbol{\theta}'$ proves the result. ∎

B. Consistency

We first have to prove that during each time slot the configurations ${}^-\boldsymbol{\theta}^{(k)}, {}^+\boldsymbol{\theta}^{(k)}$ that SDCP imposes on the cache are feasible. This is non-trivial, as the operators used in computing the allocations are defined on proper subsets of $\mathbb{R}^p$. The following lemma establishes that the allocations computed by SDCP always fall into these subsets.

Lemma 6. *The allocations $\boldsymbol{\theta}^{(k)}$ are consistent in every time slot, as they satisfy*
(a) $\boldsymbol{\theta}^{(k)} - a_k \hat{\mathbf{g}}^{(k)} \in \mathcal{C}$,
(b) $\boldsymbol{\theta}^{(k+1)} \in \mathcal{C} \cap \mathbb{R}^P_{\geq 0}$,
(c) ${}^+\boldsymbol{\theta}^{(k)}, {}^-\boldsymbol{\theta}^{(k)} \in \Theta$.

Proof: Recall that $\boldsymbol{\theta}_0 \in \mathcal{C} \cap \mathbb{R}^P_{\geq 0}$. To show (a) observe that

$$\hat{\mathbf{g}}^{(k)} \cdot \mathbf{1}_P = \sum_{j=1}^{P} \delta y_p^{(k)} \cdot D_p^{(k)} - \sum_{j=1}^{P} \delta y_p^{(k)} \cdot D_p^{(k)} = 0, \tag{16}$$

and thus if $\boldsymbol{\theta}^{(k)} \in \mathcal{C}$, then $\boldsymbol{\theta}^{(k)} - a^{(k)}\hat{\mathbf{g}}^{(k)} \in \mathcal{C}$. The definition of the Euclidean projection (10) and (a) together imply (b). Finally, observe that

$$\mathbf{1} \cdot {}^+\boldsymbol{\theta}^{(k)} = \mathbf{1} \cdot \lfloor \boldsymbol{\theta} \rfloor + \frac{P}{2} \leq \mathbf{1} \cdot \boldsymbol{\theta} + \frac{P}{2} \leq K' + \frac{P}{2} = K, \tag{17}$$

which proves (c). Note that the above motivates the choice of K' in the definition of the set of virtual allocations $\mathcal{C}$, as if $K' > K - \frac{P}{2}$ then ${}^+\boldsymbol{\theta}^{(k)}, {}^-\boldsymbol{\theta}^{(k)} \in \Theta$ may be violated due to the use of the mapping γ and $\mathbf{D}^{(k)}$ in (9). ∎

C. Convergence

To prove convergence of SDCP, we first consider the relationship between the measured miss rates ${}^+y_p^{(k)}$ and ${}^-y_p^{(k)}$ and the expected miss intensities $L_p({}^+\theta_p^{(k)})$ and $L_p({}^-\theta_p^{(k)})$, respectively. We define the measurement noise

$$\begin{aligned} {}^+\boldsymbol{\varepsilon}^{(k)} &\triangleq {}^+\vec{\mathbf{y}}^{(k)} - \vec{\mathbf{L}}({}^+\boldsymbol{\theta}^{(k)}), \\ {}^-\boldsymbol{\varepsilon}^{(k)} &\triangleq {}^-\vec{\mathbf{y}}^{(k)} - \vec{\mathbf{L}}({}^-\boldsymbol{\theta}^{(k)}) \end{aligned} \tag{18}$$

and the corresponding differences

$$\begin{aligned} \delta\boldsymbol{\varepsilon}^{(k)} &\triangleq {}^+\boldsymbol{\varepsilon}^{(k)} - {}^-\boldsymbol{\varepsilon}^{(k)}, \\ \delta\vec{\mathbf{L}}^{(k)} &\triangleq \vec{\mathbf{L}}({}^+\boldsymbol{\theta}^{(k)}) - \vec{\mathbf{L}}({}^-\boldsymbol{\theta}^{(k)}). \end{aligned} \tag{19}$$

Observe that $\mathbf{D}^{(k)}$, ${}^{+}\vec{\mathbf{y}}^{(k)}$ and ${}^{-}\vec{\mathbf{y}}^{(k)}$ are random variables and form a stochastic process. Using these definitions we can formulate the following statement about the measured miss rates.

Lemma 7. *The conditional expectation of the measurement noise and its difference satisfy*

$$\mathbb{E}[\delta\boldsymbol{\varepsilon}^{(k)}|\boldsymbol{\theta}^{(k)}] = \mathbf{0}_p. \tag{20}$$

Proof: Observe that due to the stationarity of the request arrival processes we have $\mathbb{E}[{}^{+}\boldsymbol{\varepsilon}^{(k)}|\boldsymbol{\theta}^{(k)}] = 0$ and $\mathbb{E}[{}^{-}\boldsymbol{\varepsilon}^{(k)}|\boldsymbol{\theta}^{(k)}] = 0$, which due to the additive law of expectation yields the result. ∎

Intuitively, this is equivalent to saying that the sample averages provide an unbiased estimator of the miss rates. In what follows we establish an analogous result for the update vector $\hat{\mathbf{g}}^{(k)}$ with respect to a subgradient of the interpolant $\bar{L}$ of the expected miss intensity L, which itself is a discrete function. We define and characterize $\bar{L}$ in the following lemma, which recalls known results from convex optimization.

Lemma 8. *Given the interpolants $\bar{L}_p$ of the expected miss intensities L_p of the CPs and defining the interpolant of L as $\bar{L}(\boldsymbol{\theta}) \triangleq \sum_{p=1}^{P} \bar{L}_p(\theta_p), \forall \boldsymbol{\theta} \in \mathbb{R}^P_{\geq 0}$, $\bar{L}$ is strictly convex and admits a unique minimizer $\boldsymbol{\theta}^*$ in $\mathcal{C} \cap \mathbb{R}^P_{\geq 0}$.*

Proof: Recall that each interpolant $\bar{L}_p$ of L_p is strictly convex as shown in Lemma 4. The strict convexity of $\bar{L}$ can then be obtained applying Theor. 1.17 of [21]. Then, we observe that $\boldsymbol{\theta}^*$ is the solution to a convex optimization problem with a strictly convex objective function, which is unique (Sec. 4.2.1 of [22]. ∎

For completeness, let us recall the definition of a subgradient of a function from (see, e.g., [23]).

Definition 9. Given a function $\bar{L} : \mathbb{R}^p \to \mathbb{R}$, a function $\bar{\mathbf{g}} : \mathcal{C} \subseteq \mathbb{R}^P \to \mathbb{R}^P$ is a subgradient of $\bar{L}$ over $\mathcal{C}$ iff

$$\bar{L}(\boldsymbol{\theta}') - \bar{L}(\boldsymbol{\theta}) \geq \bar{\mathbf{g}}(\boldsymbol{\theta})^T \cdot (\boldsymbol{\theta}' - \boldsymbol{\theta}), \forall \boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathcal{C}.$$

We are now ready to introduce a subgradient $\bar{\mathbf{g}}(\boldsymbol{\theta})$ for the interpolant of the expected cache miss intensity $\bar{L}$.

Lemma 10. *The function*

$$\bar{\mathbf{g}}(\boldsymbol{\theta}) \triangleq \Delta\vec{\mathbf{L}}^{(k)}(\lfloor\boldsymbol{\theta}\rfloor) - \frac{1}{P} \cdot \Delta L(\lfloor\boldsymbol{\theta}\rfloor) \cdot \mathbf{1}_P \tag{21}$$

is a subgradient of $\bar{L}$ over $\mathcal{C} \cap \mathbb{R}^P_{\geq 0}$.

Proof: Observe that for $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathcal{C}$

$$\bar{\mathbf{g}}(\boldsymbol{\theta})^T \cdot (\boldsymbol{\theta}' - \boldsymbol{\theta}) = \Delta\vec{\mathbf{L}}^{(k)}(\lfloor\boldsymbol{\theta}\rfloor)^T \cdot (\boldsymbol{\theta}' - \boldsymbol{\theta})$$
$$-\frac{1}{P} \cdot \Delta L(\lfloor\boldsymbol{\theta}\rfloor) \cdot \left[\mathbf{1}_P^T \cdot (\boldsymbol{\theta}' - \boldsymbol{\theta})\right].$$

At the same time, for $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathcal{C}$ we have

$$\mathbf{1}_P^T \cdot (\boldsymbol{\theta}' - \boldsymbol{\theta}) = (\mathbf{1}_P^T \cdot \boldsymbol{\theta}' - \mathbf{1}_P^T \cdot \boldsymbol{\theta}) = K' - K' = 0.$$

Therefore, for any $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathcal{C}$

$$\bar{\mathbf{g}}(\boldsymbol{\theta})^T \cdot (\boldsymbol{\theta}' - \boldsymbol{\theta}) = \Delta\vec{\mathbf{L}}^{(k)}(\lfloor\boldsymbol{\theta}\rfloor) \cdot (\boldsymbol{\theta}' - \boldsymbol{\theta}). \tag{22}$$

Thus, according to Def. 9, in order to show that $\bar{\mathbf{g}}$ is a subgradient of $\bar{L}$ it suffices to show that

$$\sum_{p=1}^{P} \left[\bar{L}_p(\theta'_p) - \bar{L}_p(\theta_p)\right] \geq \sum_{p=1}^{P} \Delta L_p(\lfloor\theta_j\rfloor) \cdot (\theta'_p - \theta_p). \tag{23}$$

We now show that this holds component-wise. If $\lfloor\theta'_p\rfloor - \lfloor\theta_p\rfloor = 0$, then the above clearly holds. Otherwise, if $n = \lfloor\theta'_p\rfloor - \lfloor\theta_p\rfloor \neq 0$ we apply a well known property of convex functions (Theor. 1.3.1 of [24]) to obtain:

$$\frac{\bar{L}_p(\lfloor\theta'_p\rfloor) - \bar{L}_p(\lfloor\theta_p\rfloor)}{(\lfloor\theta'_p\rfloor - \lfloor\theta_p\rfloor)} \leq \frac{\bar{L}_p(\theta'_p) - \bar{L}_p(\theta_p)}{(\theta'_p - \theta_p)}$$
$$\leq \frac{\bar{L}_p(\lfloor\theta'_p\rfloor + 1) - \bar{L}_p(\lfloor\theta_p\rfloor + 1)}{(\lfloor\theta'_p\rfloor + 1 - (\lfloor\theta_p\rfloor + 1))},$$

which, by Def. 1, can be rewritten as:

$$\frac{\Delta_n L_j(\lfloor\theta_j\rfloor)}{n} \leq \frac{\bar{L}_p(\theta'_p) - \bar{L}_p(\theta_p)}{\theta'_p - \theta_p} \leq \frac{\Delta_n L_j(\lfloor\theta_j + 1\rfloor)}{n}. \tag{24}$$

For $n > 0$ we can use the first inequality of (24) and Lemma 3 to obtain

$$\bar{L}_p(\theta'_p) - \bar{L}_p(\theta_p) \geq \Delta L_j(\lfloor\theta_j\rfloor) \cdot (\theta'_p - \theta_p). \tag{25}$$

For $n < 0$ we can use the second inequality of (24) and Lemma 3 to obtain

$$\frac{\bar{L}_p(\theta'_p) - \bar{L}_p(\theta_p)}{\theta'_p - \theta_p} \leq \Delta L_j(\lfloor\theta_j + 1\rfloor) \leq \Delta L_j(\lfloor\theta_j\rfloor) \tag{26}$$

and by multiplying the first and the second term of (26) by $\theta'_p - \theta_p$ (which is negative since $n = \lfloor\theta'_p\rfloor - \lfloor\theta_p\rfloor$ is negative), we obtain the result.

∎

The subgradient $\bar{\mathbf{g}}$ will be central to proving the convergence of SDCP, but it cannot be measured directly. The next proposition establishes a link between the update vector $\hat{\mathbf{g}}^{(k)}$, which we compute in every time slot, and the subgradient $\bar{\mathbf{g}}$.

Proposition 11. *The update vector $\hat{\mathbf{g}}^{(k)}$ is composed of the subgradient $\bar{\mathbf{g}}$ plus a component due to the noise,*

$$\hat{\mathbf{g}}^{(k)} = \bar{\mathbf{g}}(\boldsymbol{\theta}^{(k)}) + \delta\boldsymbol{\varepsilon}^{(k)} \circ \mathbf{D}^{(k)} - \frac{1}{P} \cdot \left[\delta\boldsymbol{\varepsilon}^{(k)^T} \cdot \mathbf{D}^{(k)}\right] \mathbf{1}_P.$$

Proof: We first apply (19) to obtain

$$\begin{aligned}\hat{\mathbf{g}}^{(k)} &= \delta\vec{\mathbf{L}}^{(k)} \circ \mathbf{D}^{(k)} - \frac{1}{P} \cdot (\delta\vec{\mathbf{L}}^{(k)^T} \cdot \mathbf{D}^{(k)})\mathbf{1}_P \\ &+ \delta\boldsymbol{\varepsilon}^{(k)} \circ \mathbf{D}^{(k)} - \frac{1}{P} \cdot (\delta\boldsymbol{\varepsilon}^{(k)^T} \cdot \mathbf{D}^{(k)})\mathbf{1}_P.\end{aligned} \tag{27}$$

Consider now a particular realization of the random variable $\mathbf{D}^{(k)}$. We can express component p of $\delta\vec{\mathbf{L}}^{(k)} \circ \mathbf{D}^{(k)} = \left[\vec{\mathbf{L}}({}^{+}\boldsymbol{\theta}^{(k)}) - \vec{\mathbf{L}}({}^{-}\boldsymbol{\theta}^{(k)})\right] \circ \mathbf{D}^{(k)}$ as

$$\left[L_p\left(\Pi\left(\theta_p^{(k)}\right) + \frac{1}{2}D_p^{(k)}\right) - L_p\left(\Pi\left(\theta_p^{(k)}\right) - \frac{1}{2}D_p^{(k)}\right)\right]$$
$$\cdot D_p^{(k)}$$
$$= \left[L_p\left(\Pi\left(\theta_p^{(k)}\right) + \frac{1}{2}\right) - L_p\left(\Pi\left(\theta_p^{(k)}\right) - \frac{1}{2}\right)\right]$$

$$= \left[L_p\left(\lfloor\theta_p^{(k)}\rfloor + 1\right) - L_p\left(\lfloor\theta_p^{(k)}\rfloor\right)\right],$$

where the first equality can be easily verified assuming that $D_p^{(k)} = -1$ and then assuming that it is $D_p^{(k)} = 1$. We thus obtain

$$\delta\vec{\mathbf{L}}^{(k)} \circ \mathbf{D}^{(k)} = \Delta\vec{\mathbf{L}}(\lfloor\boldsymbol{\theta}^{(k)}\rfloor)$$

and in scalar form

$$\delta\vec{\mathbf{L}}^{(k)^T} \cdot \mathbf{D}^{(k)} = \Delta L(\lfloor\boldsymbol{\theta}^{(k)}\rfloor).$$

By substituting these in (27)

$$\hat{\mathbf{g}}^{(k)} = \Delta\vec{\mathbf{L}}(\lfloor\boldsymbol{\theta}^{(k)}\rfloor) - \frac{1}{P}\cdot\Delta L(\lfloor\boldsymbol{\theta}^{(k)}\rfloor)\cdot\mathbf{1}_P$$

$$+\delta\boldsymbol{\varepsilon}^{(k)} \circ \mathbf{D}^{(k)} - \frac{1}{P}\cdot(\delta\boldsymbol{\varepsilon}^{(k)^T}\cdot\mathbf{D}^{(k)})\mathbf{1}_P$$

and using (21), we obtain the result. ∎

Furthermore, thanks to Lemma 7, the second term of (27), which is due to the noise, is zero in expectation, which provides the link between the update vector $\hat{\mathbf{g}}^{(k)}$ and the subgradient $\bar{\mathbf{g}}(\boldsymbol{\theta}^{(k)})$.

Corollary 12. *The conditional expectation of $\hat{\mathbf{g}}^{(k)}$ is $\mathbb{E}[\hat{\mathbf{g}}^{(k)}|\boldsymbol{\theta}^{(k)}] = \bar{\mathbf{g}}(\boldsymbol{\theta}^{(k)})$ and thus $\hat{\mathbf{g}}^{(k)}$ is a stochastic subgradient of $\bar{L}$, i.e. $\mathbb{E}[\hat{\mathbf{g}}^{(k)}] = \bar{\mathbf{g}}(\boldsymbol{\theta}^{(k)})$.*

This leads us to the following theorem.

Theorem 13. *The sequence $\boldsymbol{\theta}^{(k)}$ generated by SDCP converges in probability to the unique minimizer $\boldsymbol{\theta}^*$of $\bar{L}$, i.e., for arbitrary $\delta > 0$*

$$\lim_{k\to\infty} Pr\{\|\boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^*\| > \delta\} = 0.$$

Proof:

The proof of convergence is similar to (Theor. 46 in [23]), with the difference that our proof holds for Euclidean projection-based stochastic subgradients. Let us compute

$$\begin{aligned}\|\boldsymbol{\theta}^{(k+1)} - \boldsymbol{\theta}^*\|^2 &= \|\boldsymbol{\varphi}(\boldsymbol{\theta}^{(k)} - a^{(k)}\hat{\mathbf{g}}^{(k)}) - \boldsymbol{\theta}^*\|^2\\ &\leq \|\boldsymbol{\theta}^{(k)} - a^{(k)}\hat{\mathbf{g}}^{(k)} - \boldsymbol{\theta}^*\|^2\\ &= \|\boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^*\|^2 - 2a^{(k)}\cdot(\hat{\mathbf{g}}^{(k)})^T\cdot(\boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^*)\\ &\quad +(a^{(k)})^2\cdot\|\hat{\mathbf{g}}^{(k)}\|^2, \qquad (28)\end{aligned}$$

where the first inequality is due to Lemma 5. Thanks to cor 12 and Def. 9

$$\left(\mathbb{E}[\hat{\mathbf{g}}^{(k)}|\boldsymbol{\theta}^{(k)}]\right)^T\cdot(\boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^*) = \left(\bar{\mathbf{g}}(\boldsymbol{\theta}^{(k)})\right)^T\cdot(\boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^*) \geq 0.$$

Recall that the number of arriving requests per time slot $A(T)$ is bounded, and thus $\|\hat{\mathbf{g}}^{(k)}\|^2$ is bounded, i.e., $\|\hat{\mathbf{g}}^{(k)}\|^2 \leq c$ for some $0 < c < \infty$. Hence, applying the expectation to (28)

$$\mathbb{E}\left[\|\boldsymbol{\theta}^{(k+1)} - \boldsymbol{\theta}^*\|^2\middle|\boldsymbol{\theta}^{(k)}\right] \leq \|\boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^*\|^2 + c(a^{(k)})^2. \quad (29)$$

Defining the random variable

$$z_k \triangleq \|\boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^*\|^2 + c\textstyle\sum_{s=k}^{\infty}(a^{(s)})^2,$$

it can be easily verified that (29) is equivalent to the inequality $\mathbb{E}[z_{k+1}|z_k,\ldots,z_1] \leq z_k$. Consequently, $\{z_k\}_{k=1}^{\infty}$ is a supermartingale and converges almost surely to a limit z^*. Recalling now one of the required properties of the step size sequence, i.e., $\lim_{k\to\infty}\sum_{s=k}^{\infty}(a^{(k)})^2 = 0$, we have that the sequence $\{\|\boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^*\|^2\}$ also converges to z^* with probability one.

We now show by contradiction that the limit z^* is equal to zero. If this were not true, then one could find $\epsilon > 0$ and $\delta > 0$ such that, with probability $\delta > 0$, $\|\boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^*\| \geq \epsilon$ for all sufficiently large k, and thus

$$\sum_{k=0}^{\infty} a^{(k)}\cdot(\mathbb{E}[\hat{\mathbf{g}}^{(k)}|\boldsymbol{\theta}^{(k)}])^T\cdot(\boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^*) = +\infty,$$

with probability δ, which would imply

$$\mathbb{E}\left[\sum_{k=0}^{\infty} a^{(k)}\cdot\left(\bar{\mathbf{g}}(\boldsymbol{\theta}^{(k)})\right)^T\cdot(\boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^*)\right] = +\infty.$$

However, this would contradict the following relation (which is obtained by a recursion on (28) and then applying the expectation)

$$\begin{aligned}&\mathbb{E}[\|\boldsymbol{\theta}^{(k+1)} - \boldsymbol{\theta}^*\|^2] \leq\\ &\quad \|\boldsymbol{\theta}^{(0)} - \boldsymbol{\theta}^*\|^2 - 2\mathbb{E}\left[\sum_{s=0}^{k} a^{(s)}\cdot(\hat{\mathbf{g}}^{(s)})^T\cdot(\boldsymbol{\theta}^{(s)} - \boldsymbol{\theta}^*)\right] +\\ &\quad \mathbb{E}\left[\sum_{s=0}^{k} a^{(s)}\cdot\|\hat{\mathbf{g}}^{(s)})\|^2\right],\end{aligned}$$

as the left hand side cannot be negative. ∎

D. Optimality gap

It is worthwhile to note that the minimizer $\boldsymbol{\theta}^*$ of $\bar{L}$ over $\mathcal{C}\cap\mathbb{R}_{\geq 0}^P$ may not coincide with its minimizer $\boldsymbol{\theta}^{OPT}$ over Θ for two reasons: i) $K' < K$ and ii) $\boldsymbol{\theta}^{OPT}$ is forced to have integer components while $\boldsymbol{\theta}^*$ is can be a real vector. In what follows we show that the optimality gap $\|\boldsymbol{\theta}^{OPT} - \boldsymbol{\theta}^*\|_\infty$ is bounded by a small number, compared to the number of cache slots available.

Lemma 14. *The gap between the optimal solution $\boldsymbol{\theta}^{OPT}$ and the configuration $\boldsymbol{\theta}^*$ to which SDCP converges is $\|\boldsymbol{\theta}^* - \boldsymbol{\theta}^{OPT}\|_\infty \leq (3/2)P$*

Proof: We observe that $\boldsymbol{\theta}^*$ is the optimal solution of the continuous Simple Allocation Problem (SAP), expressed as $\max\left(-\sum_{p=1}^{P} L_p(\theta_p)\right)$, subject to $\sum_{p=1}^{P}\theta_p \leq K'$ with $\boldsymbol{\theta}\in\mathbb{R}_{\geq 0}^P$. K' usually referred to as *volume* and we denote with $\text{SAP}_{\text{cont}}(K')$ the problem above. The integer version of the SAP, which we denote by $\text{SAP}_{\text{int}}(K')$, is obtained from the problem above with the additional constraint $\boldsymbol{\theta}\in\mathbb{Z}^p$. According to Cor. 4.3 of [25] there exists a solution $\hat{\boldsymbol{\theta}}$ of $\text{SAP}_{\text{int}}(K')$ such that $\|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}\|_\infty \leq P$. The solution of the integer SAP can be constructed via the greedy algorithm presented in Sec. 2 of [25]. In our case, it consists of iteratively adding storage slots, one by one, each time to the CP whose miss intensity is decreased the most by using this additional slot. Based on this, it is easy to verify that a solution $\boldsymbol{\theta}^{OPT}$ can be obtained starting from $\hat{\boldsymbol{\theta}}$ and adding the remaining $K - K'$ slots. Therefore, $\|\boldsymbol{\theta}^{OPT} - \hat{\boldsymbol{\theta}}\|_\infty \leq P/2$, which implies

$$\|\boldsymbol{\theta}^{OPT} - \boldsymbol{\theta}^*\|_\infty \leq \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}\|_\infty + \|\boldsymbol{\theta}^{OPT} - \hat{\boldsymbol{\theta}}\|_\infty \leq (3/2)P.$$

∎

Algorithm 2: Conditional Step Size Sequence Computation

$a = \frac{p}{\|\hat{\mathbf{g}}^{(1)}\|_1} \cdot \frac{K'}{p}$
$b = a/10$
if $k \leq k_{BS}$; *// Bootstrap Phase*
then
  $a^{(k)} = a$
else if $k \leq M$; *// Adaptive Phase*
then
  Compute the miss-ratio $m^{(k)}$ during the current iteration
  Compute m_{5th}, i.e. the 5th percentile of the previous miss ratios $m^{(1)}, \ldots, m^{(k-1)}$
  $\hat{a}^{(k)} = a^{(k-1)}/2$
  $\tilde{a}^{(k)} = a^{(k-1)} - \frac{a^{(k-1)}-b}{M-k+1}$
  $a^{(k)} = \begin{cases} \left(\min(\hat{a}^{(k)}, \tilde{a}^{(k)}), b\right) & m^{(k)} \leq m_{5th} \\ \tilde{a}^{(k)} & \text{otherwise} \end{cases}$
else
  $a^{(k)} = a^{(k-1)} \cdot \left(1 - \frac{1}{(1+k)}\right)^{\frac{1}{2}+\epsilon}$; *// Moderate Phase*
end

V. Performance evaluation

We evaluate the performance of SDCP through simulations performed in Octave. We first describe the evaluation scenario (Sec.V-A) and show how the convergence speed is impacted by the choice of the step size sequence (Sec.V-B). We then evaluate the sensitivity of SDCP to various system parameters (Sec.V-C). Finally, recognizing that content catalogs are rarely static in the real world, we investigate the expected performance in the case of changing content catalogs (Sec.V-D).

A. Evaluation Scenario

We consider a content catalog of 10^8 objects, in line with the literature [26] and recent measurements [27]. We partition the catalog in disjoint sub-catalogs, one per each CP. We assume that the content popularity in each sub-catalog follows Zipf's law with exponent $\alpha = 0.8$, as usually done in the literature [28]. We use a cache size of $K \in \{10^4, 10^5, 10^6\}$ objects (which corresponds to cache/catalog ratios of $10^{-4}, 10^{-3}$ and 10^{-2} respectively). In practice, the request arrival rate may depend on several factors such as the cache placement in the network hierarchy, the level of aggregation, the time of day, etc. Without loss of generality, we set the request arrival rate to $\lambda = 10^2$req/s, according to recent measurements performed on ISP access networks [27]. We compare the performance of SDCP to that of the optimal allocation $\boldsymbol{\theta}^{OPT}$ (*Opt*), and to that of the naive solution in which the cache space K is equally divided among all the CPs and is unchanged throughout the simulation (*Unif*).

While we proved convergence of SDCP, the speed of convergence is crucial to let the algorithm also be of practical use: we thus consider three step size sequences, as follows. In the *Reciprocal* scheme, the step size is $a^{(k)} = a/k$, where $a = \frac{1}{\|\hat{\mathbf{g}}^{(1)}\|} \cdot \frac{K'}{p}$. Observe that, with this choice, the Euclidean norm of the first update $a^{(1)} \cdot \hat{\mathbf{g}}^{(1)}$ is $\frac{K'}{p}$, which allows to change this amount of slots in the allocation, thus obtaining a broad exploration of the allocation space at the very beginning.

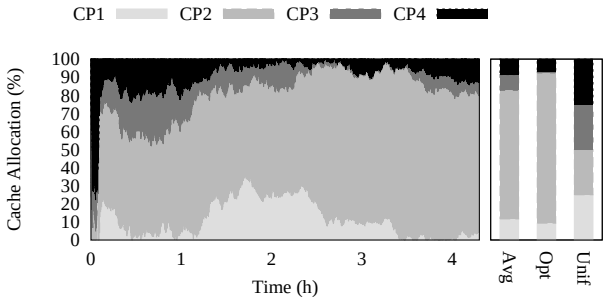

Figure 1. Evolution of the allocation of cache slots across CPs, with cache size $K = 10^5$ and *Conditional* step sizes. The three bars on the right represent the component-wise average of the allocated slots under SDCP allocation with Conditional steps (Avg), the Optimal (Opt) and Uniform (Unif) allocations.

In the *Moderate* scheme, step sizes decrease slowly, to avoid confining the exploration only to the beginning. We resort to guidelines of [29], and define the step size as $a^{(k)} = a^{(k-1)} \cdot \left(1 - \frac{1}{(1+M+k)}\right)^{\frac{1}{2}+\epsilon}$, where a is computed as above and M, ϵ are positive constants, which can be tuned to modify the decrease slope.

The third step size sequence, which we refer to as *Conditional*, is defined in Alg. 2. It consists of a *Bootstrap phase* (up to iteration k_{BS}) in which step sizes remain constant, thus allowing broad exploration. Then an *Adaptation* phase follows, up to iteration M, in which step sizes decrease, by default, linearly from an initial value a to a final value b. This decrease is steeper than linear when the miss ratio measured at the current iteration is smaller than the 5-th percentile of the miss ratio values observed so far. In this case the step size is halved, unless it already equals b. The intuition behind this phase is that we try to reduce the exploration extent every time we encounter a "good" allocation, i.e., an allocation that shows a small miss ratio compared to what we experienced so far. Note that we do not start immediately with the Adaptation phase, since we need to collect enough samples during the Bootstrap phase in order to correctly evaluate the quality of the current allocation. Finally, we continue with a *Moderate* phase, in which step sizes are updated as above and are asymptotically vanishing, thus guaranteeing convergence.

After a preliminary evaluation, we set $\epsilon = 1/100$ as in [29] and $b = a/10$. We set k_{BS} and M, i.e., the duration of the bootstrap and adaptive phases, to the number of iterations in 6 minutes and 1 hour, respectively. While the duration of these phases is clearly tied to the arrival rate, and are expected to require tuning when ported to a different scenario, we point out that performance achieved with these choices remains satisfying under the different scenarios we consider.

B. Convergence

We first consider a cache size of $K = 10^5$ and 4 CPs, receiving 13%, 75%, 2% and 10% of requests, respectively. From Fig. 1, we can observe that, after a first exploration phase, the algorithm converges to a stable allocation. It is interesting to note that the average allocation (Avg), which is obtained by averaging each component of the allocation vector throughout the iterations, is very close to the optimal one, unlike the naïve uniform allocation policy.

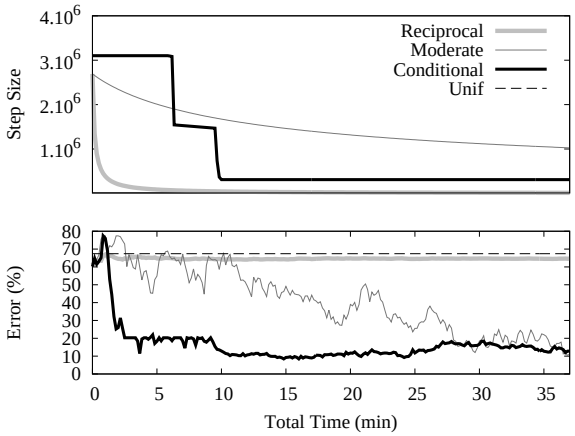

Figure 2. Error and step size sequence with cache size $K = 10^6$.

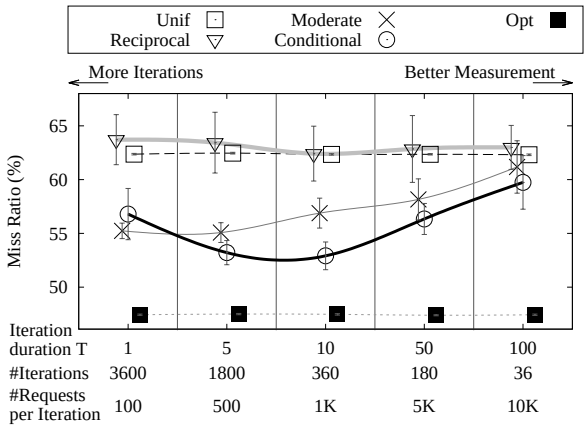

Figure 3. Impact of time-slot duration T on the average miss ratio (bars represent the 95% confidence interval over 20 runs).

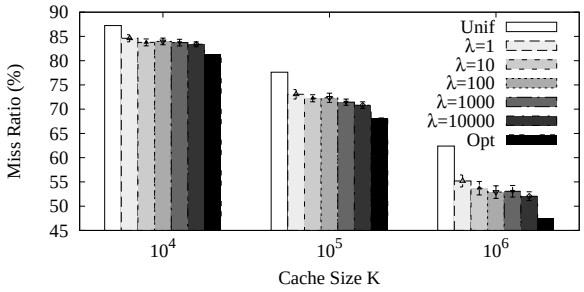

Figure 4. Miss rate measured over 1h for various average request rates λ and cache sizes K.

Second, we consider a larger scenario with cache size $K = 10^6$ and $p = 10$ CPs, one of which we is a popular CP, to which 70% of requests are directed, followed by a second one receiving 24% of requests, other 6 CPs accounting for 1% each and the remaining two CPs receiving no requests. Fig. 2 shows the step sizes and the inaccuracy of the algorithm, i.e. the distance to the optimal allocation, measured as:

$$Error(\boldsymbol{\theta}) \triangleq \frac{\|\boldsymbol{\theta}^{OPT} - \boldsymbol{\theta}\|_\infty}{K} = \frac{\max_{j=1\ldots p} |\theta_p - \theta_p^{OPT}|}{K}. \quad (30)$$

Observe that *Reciprocal* steps decrease too fast, which immediately limits the adaptation of the allocation, significantly slowing down convergence. Conversely, *Moderate* steps remain large for an overly long time, preventing the algorithm to keep the allocation in regions that guarantee good performance. *Conditional* steps show the best performance since in the Adaptation phase the step sizes are sharply decreased if the current allocation is providing a small miss ratio.

C. Sensitivity Analysis

We next study how the performance of SDCP is affected by the algorithm parameters and the scenario. We first focus on the time slot duration T. On the one hand, a small T implies that only few requests are observed in each time slot, which may result in a high noise ${}^+\boldsymbol{\varepsilon}^{(k)}, {}^-\boldsymbol{\varepsilon}^{(k)}$, and ultimately affects the accuracy of the update. On the other hand, a large T decreases the measurement noise, but allows updates to be made less frequently, which possibly slows down convergence.

To evaluate the impact of T, Fig. 3 shows the miss ratio measured over 1h for the default scenario. We consider SDCP with the three step size sequences, and compare it to the Uniform and to the Optimal allocations as benchmarks. The figure shows that SDCP with the Conditional step size sequence enhances the cache efficiency significantly. We also observe that an iteration duration of $T = 10s$ (corresponding to 100 samples on average per CP) represents a good compromise between a more accurate miss ratio estimation based on more samples (with large T) and a larger number of iterations at the cost of lower accuracy (with small T).

Fig. 4 shows the cache miss rate measured over 1h for a time slot length of $T = 10$s and for various cache sizes $K \in \{10^4, 10^5, 10^6\}$ and arrival rates $\lambda \in [1, 10^4]$. The figure confirms that the gains of SDCP hold for different cache sizes, and shows that the gain increases for large caches. To interpret the results for different arrival rates, recall that for any given time slot duration T, the average request rate affects the measurement noise. Fig. 4 confirms that the miss rate increases when the measurement noise is higher, i.e., for lower λ, but it also shows a very limited impact: the number of time slots in a relatively short time (in 1h, there are 360 time slots of duration $T = 10s$) allows SDCP to converge to a good cache configuration, in spite of the noise and the consequent estimation errors.

D. Changing Content Popularity

Recent studies [30] have observed that the catalog statistics vary over time. We show in this section that in order for SDCP to be robust to these variations, it suffices to periodically reinitialize the step sequence. To model changing content popularity, we adopt the model of [30], in which each object is characterized by a sequence of ON and OFF periods, with exponentially distributed duration T_{ON} and T_{OFF}, respectively. At each time instant, an object can be ON or OFF, and only ON objects attract requests. As in [30], we set the catalog size to $3.5 \cdot 10^6$ and the cache size to $K = 10^4$ objects. We set the average ON and OFF duration to 1 and 9 days, respectively. On average, we maintain the overall request rate of active objects equal to our default value $\lambda = 100 req/s$.

In Fig. 5 we compare *Unif* and *Conditional(τ)* that reinitialize the step sequence every τ amount of time. We consider

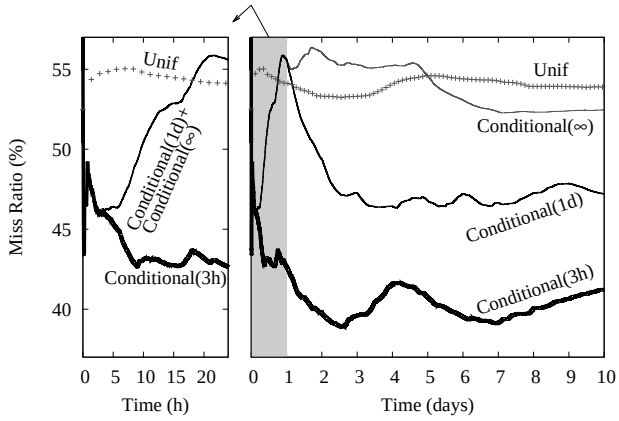

Figure 5. Miss ratio with varying content popularity. Step size sequences in *Conditional(τ)* are reset every τ interval. On the left the first day is zoomed, where *Conditional(∞)* and *Conditional($1d$)* correspond, since none of them have reinitialized the step sequence.

$\tau \in \{3h, 1d, \infty\}$, i.e., 8 reinitializations per day, daily reinitialization, or no-reinitialization, respectively. As expected, reinitialization improves cache efficiency. Indeed, already after 3 hours of simulation, the evolution of the catalog misleads *Conditional* and *Conditional(1d)* (that overlap in this time interval) causing them to have performance worse than Unif. This is expected, since *Conditional(∞)* tries to converge to the optimal allocation, which is problematic in a non-stationary scenario. At the same time, it also shows that reinitializing step sequences as in *Conditional($3h$)* is sufficient to respond to the catalog dynamics.

VI. Conclusion

One of the main challenges of in-network caching nowadays is its incompatibility with encrypted content. Our work represents a first step in solving this challenge by proposing a simple and therefore appealing system design: Stochastic Dynamic Cache Partitioning requires solely the knowledge of aggregated cache miss-intensities, based on which it provably converges to an allocation with a small optimality gap. Simulation results show the benefits of the proposed algorithm under various scenarios, and results obtained under complex content catalog dynamics further confirm the algorithm to be applicable in scenarios of high practical relevance.

Acknowledgements

This research was partially supported by Labex DigiCosme (ANR-11-LABEX-0045-DIGICOSME), part of the program "Investissement d'Avenir" Idex Paris-Saclay (ANR-11-IDEX-0003-02), and by the Swedish Foundation for Strategic Research through the Modane project. This work was carried out while Araldo was visiting KTH in the context of the EIT Digital project T15212 on "Programmable data plane for Software Defined Networks". This work benefited from support of NewNet@Paris, Cisco's Chair "Networks for the Future" at Telecom ParisTech (http://newnet.telecom-paristech.fr). Any opinion, findings or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of partners of the Chair.

References

[1] S. Deering, "Watching the waist of the protocol hourglass," *IETF51 Plenary Talk*, Aug. 2001.
[2] L. Popa, A. Ghodsi, and I. Stoica, "HTTP as the Narrow Waist of the Future Internet," in *ACM SIGCOMM HotNets Workshop*, 2010.
[3] D. Naylor, A. Finamore *et al.*, "The Cost of the "S" in HTTPS," in *ACM CoNEXT*, 2014.
[4] "IAB Statement on Internet Confidentiality," Nov. 2014.
[5] G. Barish and K. Obraczke, "World Wide Web caching: trends and techniques," *IEEE Commun. Mag.*, vol. 38, no. 5, pp. 178–184, 2000.
[6] H. S. Stone, J. Turek, and J. L. Wolf, "Optimal partitioning of cache memory," *IEEE Trans. on Comput.*, vol. 41, no. 9, pp. 1054–1068, 1992.
[7] G. Zhang, Y. Li, and T. Lin, "Caching in information centric networking: A survey," *Comput. Netw.*, vol. 57, no. 16, pp. 3128–3141, 2013.
[8] G. Carofiglio, M. Gallo *et al.*, "Evaluating per-application storage management in content-centric networks," *Comput. Commun.*, vol. 36, no. 7, pp. 750–757, 2013.
[9] A. Araldo, F. Martignon, and D. Rossi, "Representation Selection Problem : Optimizing Video Delivery through Caching," in *IFIP Networking*, 2016.
[10] S. Hoteit, M. E. Chamie *et al.*, "On Fair Network Cache Allocation to Content Providers," *Comput. Netw.*, 2015.
[11] A. Gravey, F. Guillemin, and S. Moteau, "Last Mile Caching of Video Content by an ISP," in *European Teletraffic Seminar*, 2013.
[12] D. DiPalantino and R. Johari, "Traffic Engineering vs. Content Distribution: A Game Theoretic Perspective," in *IEEE INFOCOM*, 2009.
[13] W. Jiang, R. Zhang-Shen *et al.*, "Cooperative content distribution and traffic engineering in an ISP network," *ACM SIGMETRICS Performance Evaluation Review*, vol. 37, no. 1, pp. 239–250, 2009.
[14] K. Cho, H. Jung *et al.*, "How Can an ISP Merge with a CDN?" *IEEE Commun. Mag.*, vol. 49, no. 10, pp. 156–162, 2011.
[15] B. Frank, I. Poese *et al.*, "Pushing CDN-ISP Collaboration to the Limit," *ACM SIGCOMM CCR*, vol. 43, no. 3, pp. 35–44, 2013.
[16] "Open Caching: Problem Statement and Guiding Principles," Streaming Video Alliance, Tech. Rep., 2015.
[17] Ü. Yüceer, "Discrete convexity: convexity for functions defined on discrete spaces," *Discrete Applied Mathematics*, vol. 119, no. 3, pp. 297–304, 2002.
[18] J. Carnicer and W. Dahmen, "Characterization of Local Strict Convexity Preserving Interpolation Methods by C1 Functions," *J. of Approximation Theory*, vol. 77, no. 1, pp. 2–30, 1994.
[19] W. Wang and M. A. Carreira-Pepinan, "Projection onto the probability simplex : An efficient algorithm with a simple proof and an application," *arXiv:1309.1541v1*, 2013.
[20] H. H. Bauschke and J. M. Borwein, "On Projection Algorithms for Solving Convex Feasibility Problems," *SIAM Review*, vol. 38, no. 3, pp. 367–426, 1996.
[21] S. M. Sefanov, *Separable Programming: Theory and Methods*. Springer, 2013.
[22] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
[23] N. Z. Shor, *Nondifferentiable Optimization and Polynomial Problems*. Springer, 1998.
[24] C. P. Niculescu and L.-E. Persson, *Convex Functions and Their Applications: A Contemporary Approach*. Springer, 2004.
[25] D. S. Hochbaum, "Lower and Upper Bounds for the Allocation Problem and Other Nonlinear Optimization Problems," *Mathematics of Operation Research*, vol. 19, no. 2, pp. 390–409, 1994.
[26] C. Fricker, P. Robert *et al.*, "Impact of Traffic Mix on Caching Performance in a Content-Centric Network," in *IEEE INFOCOM*, 2012.
[27] C. Imbrenda, L. Muscariello, and D. Rossi, "Analyzing Cacheable Traffic in ISP Access Networks for Micro CDN Applications via Content-Centric Networking," in *ACM SIGCOMM ICN*, 2014.
[28] K. Poularakis, G. Iosifidis, and L. Tassiulas, "Joint Caching and Base Station Activation for Green Heterogeneous Cellular Networks," in *IEEE ICC*, 2015.
[29] Q. Wang, "Optimization with Discrete Simultaneous Perturbation Stochastic Approximation Using Noisy Loss Function Measurement," Ph.D. dissertation, John Hopkins University, 2013.
[30] M. Garetto, E. Leonardi, and S. Traverso, "Efficient analysis of caching strategies under dynamic content popularity," *IEEE INFOCOM*, 2015.